

Juniper Apstra Drain Mode Guide

Published
2024-10-16

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Juniper Apstra Drain Mode Guide

Copyright © 2024 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <https://support.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About This Guide | iv

1

Introduction

Introduction | 2

2

Activate or Disable Drain Mode

Activate or Disable Drain Mode | 4

3

IBA Monitoring of Devices in Drain Mode

IBA Monitoring of Devices in Drain Mode | 7

4

Configuration Examples

Configuration Examples | 10

Drain Spine Devices (L2 and L3 Blueprints) | 10

Drain Leaf Devices (Server-Facing Ports w/ MLAG) | 13

Drain Leaf Devices (L2 Server-Facing Ports no MLAG) | 18

Drain Leaf Devices (L3 Connected Servers) | 23

About This Guide

This guide provides information about using Drain Mode in Juniper Apstra, with configuration examples. Drain Mode enables you to gracefully drain traffic from devices without shutting down the BGP neighbor routes.

1


CHAPTER

Introduction

Introduction | 2


Introduction

Juniper Apstra supports Drain Mode for managed switches, allowing the operator to gracefully drain traffic from devices without simply shutting down the BGP neighbor relationships. This is implemented through modifications to the BGP process (inbound/outbound route-maps), shutting down connected L2 server ports, and shutting down MLAG peer link ports. By using Drain Mode, operators can minimize the number of dropped/lost traffic during these operations. During maintenance, redundancy is handled by ECMP/MLAG as long as there are suitable redundant systems in place. A visual example of Drain Mode on Spine switches is displayed below:




Normal Operation

L1 has two paths to L2 with ECMP
 $L1 \rightarrow S1 \rightarrow L2$
 • $L1 \rightarrow S2 \rightarrow L2$
 Traffic is evenly distributed across both paths automatically



Entering Maintenance Mode

Route-map is added to S1
 Filtering of outbound and inbound routes
 Route through S1 is removed from tables
 Only one remaining route exists
 • $L1 \rightarrow S2 \rightarrow L2$
 TCP flows are moved to the remaining route
 S1 is "drained" of traffic



Maintenance Mode

S1 is drained of traffic and can be rebooted, upgraded, downgraded without impact to network traffic

2

CHAPTER

Activate or Disable Drain Mode

Activate or Disable Drain Mode | 4


Activate or Disable Drain Mode

IN THIS SECTION


- [Activate Drain Mode | 4](#)
- [Disable Drain Mode | 5](#)

Activate Drain Mode

Activate Drain Mode by switching devices to the Drain state in Juniper Apstra:



Once the device is switched to Drain, the change must be completed with the Commit button. The following image shows an example workflow using the Drain functionality.



Disable Drain Mode

To restore a device to service, switch the Deploy Mode setting back to Deploy, then Commit.

3

CHAPTER

IBA Monitoring of Devices in Drain Mode

IBA Monitoring of Devices in Drain Mode | 7

IBA Monitoring of Devices in Drain Mode

IN THIS SECTION

- [Example | 7](#)
- [Recommended Usage | 8](#)

A prebuilt IBA (intent-based analytics) probe is available in Juniper Apstra. You can activate it by instantiating a predefined probe named “Drain traffic anomaly”. The required value for **Threshold in bps** works as follows:

- Value is the net sum of traffic on all hosted_interfaces
- This does not include traffic on the Ethernet management port which is not part of the probe measurement
- These interfaces include all L3 BGP enabled paths
- Server facing interfaces are shut during Drain Mode and are not part of this calculation
- The threshold describes the amount of traffic you wish to be alerted on (above the value) if devices are in the Drain state
- This ensures that you do not perform actual maintenance operations on a device that has not been fully drained.

Example

Spine1 is connected to 4 leaf switches, each connection runs the eBGP routing process. All application (server) based traffic flows are rehashed via ECMP onto other links and the basic BGP neighbor updates are still running. In a lab example with a small topology, this is effectively 1.5KBPS per link. With 4 neighbors, the total traffic we expect to remain on the devices is approximately 6KBPS. If we set the probe **Threshold in bps** to 10KBPS (10000), the probe generates anomalies if there is more than 10K on all of the 4 interfaces combined.

Recommended Usage

Enable the probe with 100KBPS and leave it running in all Blueprints. When a device enters the Drain state, an anomaly appears as the traffic is removed from the links. This anomaly should only exist for a few seconds. If the anomaly does not clear, the device is not fully in Drain Mode. Once the anomaly clears, you are free to switch the device to the Ready state to take it out of service completely. It is also possible that you will not see the anomaly as it may appear and disappear very quickly.

4

CHAPTER

Configuration Examples

Configuration Examples | 10

Configuration Examples

IN THIS SECTION

- [Drain Spine Devices \(L2 and L3 Blueprints\) | 10](#)
- [Drain Leaf Devices \(Server-Facing Ports w/ MLAG\) | 13](#)
- [Drain Leaf Devices \(L2 Server-Facing Ports no MLAG\) | 18](#)
- [Drain Leaf Devices \(L3 Connected Servers\) | 23](#)

The following sections provide Drain Mode configuration examples for different OS and device combinations.

Drain Spine Devices (L2 and L3 Blueprints)


IN THIS SECTION

- [Drain \(NX-OS\) | 11](#)
- [Drain \(Junos\) | 12](#)

The following occurs when draining the Spine:

- Outbound routes are removed from the device's routing table.
- Routes to destinations with the device's ASN (Autonomous System Numbers) in the AS-PATH are removed from all devices in the network.
- Packets are forwarded through remaining ECMP (Equal Cost Multi-Path) paths for all destinations.

NOTE: It is highly unlikely that a single in-flight packet will be lost. This is dependent however, on the L3 ECMP to L2 path hashing algorithms in the hardware and NOS.



Drain (NX-OS)

```

ip prefix-list Drain seq 5 permit 0.0.0.0/0 le 32
route-map Drain deny 10
  match ip address prefix-list Drain
  exit
!
neighbor 172.16.0.1 remote-as 64514
  address-family ipv4 unicast
    route-map Drain out
    route-map Drain in
  exit
exit
neighbor 172.16.0.3 remote-as 64514
  address-family ipv4 unicast
    route-map Drain out
    route-map Drain in
  exit
exit

```

Drain (Junos)

```
[edit policy-options]
+ route-filter-list Drain {
+   0.0.0.0/0 upto /32;
+ }
[edit policy-options]
+ policy-statement Drain {
+   term Drain-10 {
+     from {
+       family inet;
+       route-filter-list Drain;
+     }
+     then reject;
+   }
+ }
[edit protocols bgp group l3clos-s neighbor 172.16.0.7]
+ import ( Drain );
- export ( SPINE_TO_LEAF_FABRIC_OUT && BGP-AOS-Policy );
+ export ( Drain );
[edit protocols bgp group l3clos-s neighbor 172.16.0.9]
+ import ( Drain );
- export ( SPINE_TO_LEAF_FABRIC_OUT && BGP-AOS-Policy );
+ export ( Drain );
[edit protocols bgp group l3clos-s neighbor 172.16.0.11]
+ import ( Drain );
- export ( SPINE_TO_LEAF_FABRIC_OUT && BGP-AOS-Policy );
+ export ( Drain );
[edit protocols bgp group l3clos-s-evpn neighbor 10.0.0.0]
+ import ( Drain );
- export ( SPINE_TO_LEAF_EVPN_OUT );
+ export ( Drain );
[edit protocols bgp group l3clos-s-evpn neighbor 10.0.0.1]
+ import ( Drain );
- export ( SPINE_TO_LEAF_EVPN_OUT );
+ export ( Drain );
[edit protocols bgp group l3clos-s-evpn neighbor 10.0.0.2]
+ import ( Drain );
- export ( SPINE_TO_LEAF_EVPN_OUT );
+ export ( Drain );
```

Drain Leaf Devices (Server-Facing Ports w/ MLAG)

IN THIS SECTION

- [Drain \(NX-OS\) | 14](#)
- [Drain \(EOS\) | 15](#)
- [Undrain \(NS-OS\) | 16](#)
- [Undrain \(EOS\) | 17](#)

The following occurs when draining Leaf devices with a server-facing port in an MLAG:

- A route-map is placed on all BGP neighbors restricting inbound and outbound routes.
- Server facing interfaces are shutdown.
- MLAG peer interfaces are shutdown.

What happens at L3:


- Outbound routes are removed from the device's routing table.
- Routes to destinations with the device's ASN in the AS-PATH are removed from all devices in the network.
- Packets are forwarded through remaining ECMP paths for all destinations.

NOTE: It is highly unlikely that a single in-flight packet will be lost, however, this is dependent on the L3 ECMP to L2 path hashing algorithms in the hardware and NOS.

What happens at L2:

- Server interfaces to this device will go DOWN.
- Packets from the server that happen to be hashed onto this device via MLAG may be dropped depending on where they are in the forwarding process.
- Packets from the server that happen to be hashed onto this device via MLAG may be forwarded over the MLAG peer link depending on where they are in the forwarding process.
- Flows will be reestablished on the alternate MLAG interfaces.

- New flows will be established on the remaining MLAG interfaces.



Drain (NX-OS)

```

interface Ethernet1/1
shutdown
exit
!
interface Ethernet1/2
shutdown
exit
!
interface port-channel1
shutdown
exit
!
ip prefix-list Drain seq 5 permit 0.0.0.0/0 le 32
ipv6 prefix-list DrainV6 seq 5 permit 0::0/0 le 128
route-map Drain deny 10
  match ip address prefix-list Drain
  exit
!
route-map DrainV6 deny 10
  match ipv6 address prefix-list DrainV6
  exit
!

```

```

router bgp 64514
  neighbor 10.0.0.0 remote-as 64512
  address-family l2vpn evpn
    route-map Drain out
    route-map Drain in
  exit
exit
neighbor 172.16.0.0 remote-as 64512
address-family ipv4 unicast
  route-map Drain out
  route-map Drain in
  exit
exit

```

Drain (EOS)

```

interface Ethernet5
  shutdown
  exit
!
interface Ethernet6
  shutdown
  exit
!
interface port-channel1
  shutdown
  exit
!
interface port-channel2
  shutdown
  exit
!
ip prefix-list Drain seq 5 permit 0.0.0.0/0 le 32
route-map Drain deny 10
  match ip address prefix-list Drain
  exit
!
router bgp 102
  neighbor 10.10.4.0 route-map Drain out
  neighbor 10.10.4.0 route-map Drain in
  neighbor 10.10.4.8 route-map Drain out

```

```

neighbor 10.10.4.8 route-map Drain in
default neighbor 10.10.4.19 route-map MlagPeer out
neighbor 10.10.4.19 route-map Drain out
neighbor 10.10.4.19 route-map Drain in
!
```

Undrain (NS-OS)

What happens at L2:

- Server interface to this device will go UP
- New flows will be hashed onto the newly available MLAG interface

```

interface Ethernet1/1
  no shutdown
  exit
!
interface Ethernet1/2
  no shutdown
  exit
!
interface port-channel1
  no shutdown
  exit
!
no ip prefix-list Drain seq 5 permit 0.0.0.0/0 le 32
no ipv6 prefix-list DrainV6 seq 5 permit 0::0/0 le 128
no route-map Drain deny 10
!
no route-map DrainV6 deny 10
!
router bgp 64514
  neighbor 10.0.0.0 remote-as 64512
  address-family l2vpn evpn
    default route-map Drain out
    default route-map Drain in
  exit
exit
```

Undrain (EOS)

What happens at L2:

- Server interface to this device will go UP
- New flows will be hashed onto the newly available MLAG interface

```
interface Ethernet5
  no shutdown
  exit
!
interface Ethernet6
  no shutdown
  exit
!
interface port-channel1
  no shutdown
  exit
!
interface port-channel2
  no shutdown
  exit
!
no ip prefix-list Drain seq 5 permit 0.0.0.0/0 le 32
no route-map Drain deny 10
!
router bgp 102
  default neighbor 10.10.4.0 route-map Drain out
  default neighbor 10.10.4.0 route-map Drain in
  default neighbor 10.10.4.8 route-map Drain out
  default neighbor 10.10.4.8 route-map Drain in
  default neighbor 10.10.4.19 route-map Drain out
  neighbor 10.10.4.19 route-map MlagPeer out
  default neighbor 10.10.4.19 route-map Drain in
!
```

Drain Leaf Devices (L2 Server-Facing Ports no MLAG)

IN THIS SECTION

- [Drain \(Junos\) | 18](#)
- [Drain \(NX-OS\) | 20](#)
- [Drain \(EOS\) | 20](#)
- [Undrain \(NX-OS\) | 21](#)
- [Undrain \(EOS\) | 22](#)

The following occurs when draining a Leaf device with a server-facing port with no MLAG:

- A route-map is placed on all BGP neighbors restricting inbound and outbound routes
- Server facing interfaces are shutdown

Drain (Junos)

```
[interfaces replace: ae1]
+ disable;

[interfaces replace: xe-0/0/2]
+ disable;

[interfaces replace: xe-0/0/3]
+ disable;

[routing-instances blue protocols bgp group l3rtr neighbor 192.168.0.11]
- import ( RoutesFromExt-blue-Default_immutable );
- export ( RoutesToExt-blue-Default_immutable );
+ import ( Drain );
+ export ( Drain );

[routing-instances red protocols bgp group l3rtr neighbor 192.168.0.7]
- import ( RoutesFromExt-red-Default_immutable );
- export ( RoutesToExt-red-Default_immutable );
+ import ( Drain );
```

```

+ export ( Drain );

[protocols bgp group l3clos-1 neighbor 172.16.0.2]
- export ( LEAF_TO_SPINE_FABRIC_OUT && BGP-AOS-Policy );
+ import ( Drain );
+ export ( Drain );

[protocols bgp group l3clos-1 neighbor 172.16.0.8]
- export ( LEAF_TO_SPINE_FABRIC_OUT && BGP-AOS-Policy );
+ import ( Drain );
+ export ( Drain );

[protocols bgp group l3clos-1-evpn neighbor 10.0.0.3]
- export ( LEAF_TO_SPINE_EVPN_OUT && EVPN_EXPORT );
+ import ( Drain );
+ export ( Drain && EVPN_EXPORT );

[protocols bgp group l3clos-1-evpn neighbor 10.0.0.4]
- export ( LEAF_TO_SPINE_EVPN_OUT && EVPN_EXPORT );
+ import ( Drain );
+ export ( Drain && EVPN_EXPORT );

[protocols bgp group l3rtr neighbor 192.168.0.3]
- import ( RoutesFromExt-default-Default_immutable );
- export ( RoutesToExt-default-Default_immutable );
+ import ( Drain );
+ export ( Drain );

+ [policy-options route-filter-list Drain]
+ 0.0.0.0/0 upto /32;

+ [policy-options policy-statement Drain term Drain-10 from]
+ route-filter-list Drain;
+ family inet;

+ [policy-options policy-statement Drain term Drain-10]
+ then reject

```

Drain (NX-OS)

```
interface Ethernet1/41
  shutdown
  exit
!
ip prefix-list Drain seq 5 permit 0.0.0.0/0 le 32
route-map Drain deny 10
  match ip address prefix-list Drain
  exit
!
router bgp 64516
  neighbor 172.16.0.8 remote-as 64512
  address-family ipv4 unicast
    route-map Drain out
    route-map Drain in
  exit
  exit
  neighbor 172.16.0.22 remote-as 64513
  address-family ipv4 unicast
    route-map Drain out
    route-map Drain in
  exit
  exit
exit
!
```

Drain (EOS)

```
interface Ethernet5
  shutdown
  exit
!
ip prefix-list Drain seq 5 permit 0.0.0.0/0 le 32
route-map Drain deny 10
  match ip address prefix-list Drain
  exit
!
router bgp 104
  default neighbor 9.0.0.1 route-map RoutesToExt out
```

```

neighbor 9.0.0.1 route-map Drain out
default neighbor 9.0.0.1 route-map RoutesFromExt in
neighbor 9.0.0.1 route-map Drain in
neighbor 10.10.4.4 route-map Drain out
neighbor 10.10.4.4 route-map Drain in
neighbor 10.20.30.4 route-map Drain out
neighbor 10.20.30.4 route-map Drain in
neighbor 10.10.4.12 route-map Drain out
neighbor 10.10.4.12 route-map Drain in
neighbor 10.20.30.5 route-map Drain out
neighbor 10.20.30.5 route-map Drain in
vrf Finance
    default neighbor 9.0.0.1 route-map RoutesToExt-Finance out
    neighbor 9.0.0.1 route-map Drain out
    default neighbor 9.0.0.1 route-map RoutesFromExt-Finance in
    neighbor 9.0.0.1 route-map Drain in
exit
!
```

Undrain (NX-OS)

```

interface Ethernet1/41
    no shutdown
    exit
!
no ip prefix-list Drain seq 5 permit 0.0.0.0/0 le 32
no route-map Drain deny 10
!
router bgp 64516
    neighbor 172.16.0.8 remote-as 64512
    address-family ipv4 unicast
        default route-map Drain out
        default route-map Drain in
    exit
    exit
neighbor 172.16.0.10 remote-as 64512
    address-family ipv4 unicast
        default route-map Drain out
        default route-map Drain in
    exit
exit
```

```

neighbor 10.0.0.1 remote-as 64513
  address-family l2vpn evpn
    default route-map Drain out
    default route-map Drain in
  exit
exit
neighbor 172.16.0.20 remote-as 64513
  address-family ipv4 unicast
    default route-map Drain out
    default route-map Drain in
  exit
exit
neighbor 172.16.0.22 remote-as 64513
  address-family ipv4 unicast
    default route-map Drain out
    default route-map Drain in
  exit
exit
exit
!
```

Undrain (EOS)

```

interface Ethernet5
  no shutdown
  exit
!
no ip prefix-list Drain seq 5 permit 0.0.0.0/0 le 32
no route-map Drain deny 10
!
router bgp 104
  default neighbor 9.0.0.1 route-map Drain out
  neighbor 9.0.0.1 route-map RoutesToExt out
  default neighbor 9.0.0.1 route-map Drain in
  neighbor 9.0.0.1 route-map RoutesFromExt in
  default neighbor 10.10.4.4 route-map Drain out
  default neighbor 10.10.4.4 route-map Drain in
  default neighbor 10.20.30.4 route-map Drain out
  default neighbor 10.20.30.4 route-map Drain in
  default neighbor 10.10.4.12 route-map Drain out
  default neighbor 10.10.4.12 route-map Drain in
```

```


default neighbor 10.20.30.5 route-map Drain out
default neighbor 10.20.30.5 route-map Drain in
vrf Finance
  default neighbor 9.0.0.1 route-map Drain out
  neighbor 9.0.0.1 route-map RoutesToExt-Finance out
  default neighbor 9.0.0.1 route-map Drain in
  neighbor 9.0.0.1 route-map RoutesFromExt-Finance in
exit
!
```

Drain Leaf Devices (L3 Connected Servers)

IN THIS SECTION

- Drain (EOS) | 24
- Undrain (EOS) | 24

The following occurs when draining a Leaf device with a server connected at L3.



Drain (EOS)

```
ip prefix-list Drain seq 5 permit 0.0.0.0/0 le 32
route-map Drain deny 10
  match ip address prefix-list Drain
  exit
!
router bgp 102
  neighbor 10.10.4.0 route-map Drain out
  neighbor 10.10.4.0 route-map Drain in
  neighbor 10.10.4.8 route-map Drain out
  neighbor 10.10.4.8 route-map Drain in
  neighbor 11.0.0.1 route-map Drain out
  neighbor 11.0.0.1 route-map Drain in
!
```

Undrain (EOS)

```
no ip prefix-list Drain seq 5 permit 0.0.0.0/0 le 32
no route-map Drain deny 10
!
router bgp 102
  default neighbor 10.10.4.0 route-map Drain out
  default neighbor 10.10.4.0 route-map Drain in
  default neighbor 10.10.4.8 route-map Drain out
  default neighbor 10.10.4.8 route-map Drain in
  default neighbor 11.0.0.1 route-map Drain out
  default neighbor 11.0.0.1 route-map Drain in
!
```

SEE ALSO

| [Drain Device Traffic](#)